
Robustifying Sequential Neural Processes

Jaesik Yoon¹ Gautam Singh² Sungjin Ahn^{2,3}

Abstract

When tasks change over time, *meta-transfer learning* seeks to improve the efficiency of learning a new task via both *meta-learning* and *transfer-learning*. While the standard attention has been effective in a variety of settings, we question its effectiveness in improving meta-transfer learning since the tasks being learned are dynamic and the amount of context can be substantially smaller. In this paper, using a recently proposed meta-transfer learning model, Sequential Neural Processes (SNP), we first empirically show that it suffers from a similar underfitting problem observed in the functions inferred by Neural Processes. However, we further demonstrate that unlike the meta-learning setting, the standard attention mechanisms are not effective in meta-transfer setting. To resolve, we propose a new attention mechanism, Recurrent Memory Reconstruction (RMR), and demonstrate that providing an imaginary context that is recurrently updated and reconstructed with interaction is crucial in achieving effective attention for meta-transfer learning. Furthermore, incorporating RMR into SNP, we propose Attentive Sequential Neural Processes-RMR (ASNP-RMR) and demonstrate in various tasks that ASNP-RMR significantly outperforms the baselines.

1. Introduction

A central challenge in machine learning is to improve learning efficiency. Among such approaches are meta-learning (Schmidhuber, 1987; Bengio et al., 1990) and transfer learning (Pratt, 1993; Pan & Yang, 2009). Meta-learning aims to learn the learning process itself and thus enables efficient learning (e.g., from a small amount of data) while

¹SAP ²Department of Computer Science, Rutgers University ³Rutgers Center for Cognitive Science. Correspondence to: Jaesik Yoon <jaesik.yoon.kr@gmail.com>, Sungjin Ahn <sjn.ahn@gmail.com>.

transfer learning allows efficient warm-start of a new task by transferring knowledge from previously learned tasks.

In many scenarios, these two problems are not separate but appear in a combined way. For example, to build a customer preference model monthly, we would like to build it by transferring knowledge from the models of the previous months instead of starting from scratch, because the general preference of a customer would not change much across months. However, due to the monthly preference shift, e.g., due to a seasonal change, we also need to efficiently learn from the new observations of the new month.

Sequential Neural Processes (SNP) (Singh et al., 2019) is a new probabilistic model class to resolve the above-mentioned *meta-transfer learning* problem. In SNP, meta-transfer learning is modeled as stochastic processes that change with time (thus, a stochastic process of stochastic processes). At each time-step, the stochastic process—or equivalently a *task*—is meta-learned from the context. SNP represents each task by a latent state of the Neural Process (NP) and models the temporal dynamics of such latent states using a recurrent state-space model (Hafner et al., 2018). This enables SNP to transfer the knowledge of the previous tasks to learn a new task.

It is well-known that Neural Processes (NP) suffers from the *underfitting* problem because all context observations are encoded with limited expressiveness (e.g., by sum or mean encoding) to satisfy the order-invariant property (Wagstaff et al., 2019). In Kim et al. (2019), the authors observe that query-dependent attention can substantially mitigate the problem and propose Attentive Neural Processes (ANP). Therefore, a crucial next question is whether SNP—which is partly based on NP for task-level meta-learning but also equipped with temporal-transfer—also suffers from underfitting, and if so, how we can resolve the problem.

In this paper, we argue not only that there is underfitting in SNP but also that it affects the robustness more severely in comparison to NP. We observe that this is because of two novel problems, *sparse context* and *obsolete context*, that occur in the novel setting of meta-transfer learning. In Singh et al. (2019), it is shown that in comparison to meta-learning, meta-transfer learning is expected to learn a task more efficiently by using a much smaller amount of context or even empty context due to the availability of temporal transfer. However, this sparsity becomes an issue because




Figure 1: Task shift. A context observation $c_1 = (x_1, y_1)$ is provided for a task T_1 (black line). Then, the task is changed to T_2 (blue line). After the task is shifted to T_2 , the value $f(x_2)$ is queried whose true value is y_2 . A standard attention will use the obsolete context $c_1 = (x_1, y_1)$ to infer $f(x_2)$ and a high attention-weight will be given to it because x_1 and x_2 are close. As a result, the attention will suggest a high value for $f(x_2)$ even though the true value y_2 is small. Our proposed model reconstructs c_1 so that its value can be properly adapted to the new task (T_2).

when the context is small or empty, the attention—which is a remedy for underfitting—becomes highly ineffective or even not applicable. In the case of sparse context, it is possible to include the past contexts also for attention. However, this raises the second problem, the problem of obsolete context, because the past contexts come from tasks that are different from the current task. Thus, we argue that if the past contexts are not properly transformed for the current task, attention on them may hurt the performance. An illustration is shown in Fig. 1.

To this end, we propose a novel attention mechanism for meta-transfer learning, called Recurrent Memory Reconstruction and Attention (RMRA). In RMRA, to resolve the sparsity problem, for each task, we augment its original context with a generated *imaginary* context. Thus, even if a task provides a small or empty context, we can still apply attention effectively on the generated imaginary context. To resolve the obsolete context problem, we generate this imaginary context using a novel *Recurrent Memory Reconstruction* (RMR) mechanism. RMR temporally encodes *all* the past context observations using recurrent updates and then reconstructs a reformed imaginary context for each task. In this way, the past contexts are properly transformed into a useful representation for the current task. In addition, we do not need to limit the attention to a finite-size window but can use the entire past information without storing it explicitly. By augmenting SNP with RMRA, we propose a novel and robust SNP model, called Attentive SNP-RMR.

Our main contributions are: (i) we identify why SNP should also suffer from underfitting and empirically show that it is indeed the case, (ii) we provide reasons why the existing attention mechanisms for NP are sub-optimal in the meta-transfer learning setting — due to sparse and obsolete context — and provide empirical analysis for it, (iii) we propose a novel Recurrent Memory Reconstruction and Attention (RMRA) mechanism to resolve the problem and we combine RMRA with SNP and thereby propose a novel

model *Attentive SNP-RMR*, and (iv) in our experiments, we empirically show that the RMRA mechanism resolves underfitting efficiently and effectively and results in superior performance in comparison to the baselines in various meta-transfer learning settings.

2. Background

Neural Process (NP) (Garnelo et al., 2018) learns to learn a task τ to map an input $x \in \mathbb{R}^{d_x}$ to an output $y \in \mathbb{R}^{d_y}$ given a *context* dataset $C = (X_C, Y_C) = \{(x^{(n)}, y^{(n)})\}_{n \in [N_C]}$. Here, N_C is the number of data points in C and $[N_C] \equiv \{1, \dots, N_C\}$. To learn a task distribution from this context, NP uses a distribution $P(z|C)$ to sample a task representation z . This makes NP a probabilistic meta-learning framework. Next, an observation model $p(y|x, z)$ takes an input x and returns an output y . The generative process for NP conditioned on the context is given by:

$$P(Y, z|X, C) = P(Y|X, z)P(z|C) \quad (1)$$

where $P(Y|X, z) = \prod_{n \in [N_D]} P(y^{(n)}|x^{(n)}, z)$ and $D = (X, Y) = \{(x^{(n)}, y^{(n)})\}_{n \in [N_D]}$ is the *target* dataset. To obtain the training data for this meta-learning setting, we draw multiple tasks from the true task distribution and sample (C, D) for each task. Note that to implement NP, C is encoded via a permutation-invariant function, such as $\sum_n \text{MLP}(x^{(n)}, y^{(n)})$. Kim et al. (2019) argue that such a sum-aggregation produces an encoding that is not expressive enough and consequently hurts the observation model $P(Y|X, z)$. This is a key limitation of NP and is addressed by Attentive Neural Processes (ANP).

Attentive Neural Process (ANP) Kim et al. (2019) identify the problem of *underfitting* in NP. In underfitting, tasks learned from the context set fail to accurately predict the target points, including those in the context set. Also, the learned task distribution shows high uncertainty. Authors show that using a larger latent z is also not sufficient. Therefore to address this, ANP combines the observation model $p(y|x, z)$ with attention on the context points. This allows the model to take a query x and perform attention on the most relevant data points to predict the target output y . To achieve this, an attention function $\text{Attend}(C; x^q)$ is implemented, which takes the context C and a query x^q and returns a read value r_{x^q} .

$$r_{x^q} = \text{Attend}(C; x^q) = \sum_{n \in [N_C]} w_n y^{(n)} \quad (2)$$

Here, for each $y^{(n)} \in Y_C$, there is an associated weight w_n which is computed using a similarity function, $\text{sim}(X_C, x^q)$. Using r_x , the observation model becomes $P(y|x, z, r_x)$ and results in the following generative process.

$$P(y, z|x, C) = P(y|x, z, r_x)P(z|C) \quad (3)$$

where $(x, y) \in D$ is a target point. The ANP framework allows for the use of a variety of attentive mechanisms (Vaswani et al., 2017).

Sequential Neural Process (SNP) (Singh et al., 2019) While the goal of meta-learning is to learn a single task distribution from a given context C , many situations consist of a sequence of tasks $\{\tau_t\}_{t \in [T]}$ that change gradually with time. Here, t is a task-step. Thus, a framework that learns a task τ_t from a context C_t must also utilize its similarity with the previous tasks to be able to learn with less context. Let z_t denote a representation for a task τ_t . Then SNP uses a distribution $p(z_t | z_{<t}, C_t)$ to learn from both C_t and the previous task representations $z_{<t}$. From this standpoint, SNP is a meta-transfer learning framework. Given a task representation z_t , the model takes a query x_t and generates an output y_t through an observation model $p(y_t | x_t, z_t)$.

For a task τ_t , let $C_t = \{(x_t^{(n)}, y_t^{(n)})\}$ be the context and let (x_t, y_t) be a target point. Then, we describe the generative process for the target conditioned on the context as follows.

$$P(Y, Z | X, C) = \prod_{t=1}^T P(y_t | x_t, z_t) P(z_t | z_{<t}, C_t). \quad (4)$$

Here, C , X , Y and Z respectively denote the set aggregations of C_t , x_t , y_t and z_t over the entire roll-out.

3. Attentive Meta-Transfer

In this section, we describe our proposed model to resolve the problem of underfitting using attention in the setting of meta-transfer learning. We first propose a novel recurrent attention mechanism called Recurrent Memory Reconstruction (RMR), which resolves the problem of *sparse context* that makes underfitting in SNP more severe and *obsolete context* which turns the past contexts to noise due to task shift. We then propose to incorporate RMR into the SNP framework to obtain a robust meta-transfer learning model, Attentive Sequential Neural Processes-RMR (ASNP-RMR).

3.1. Recurrent Memory Reconstruction

To resolve the limitations of the standard attention models in dealing with sparse and obsolete context, we propose the Recurrent Memory Reconstruction (RMR) mechanism. The key ideas in RMR are (i) to generate *imaginary context* to complement the sparse context and (ii) to introduce *recurrent reformation* to appropriately transform the obsolete context to a new useful representation upon a task-shift. The imaginary context is generated at each task-step using the updated representation of past contexts via recurrent reformation.

The main task of RMR at each task-step t is to generate a new imaginary context \tilde{C}_t from an encoded representation of the past contexts (both real and imaginary). The imaginary context contains K memory cells, each of which is a

Algorithm 1 Recurrent Memory Reconstruction

ensure: $f_{xy}^c = \text{MLP}$, $f_{\text{order-invariant}} = \sum f_{xy}^c$
initialize: $h_0^x, \{\tilde{h}_0^k\}_{k \in [K]}$, $\tilde{X}_0, \{\tilde{v}_0^k\}_{k \in [K]}$
for $t \in [T]$ **do**
 # Context Processing
 $r_t^c \leftarrow f_{\text{order-invariant}}(C_t)$
 $C_t \leftarrow \{(x_t^{(n)}, f_{xy}^c(x_t^{(n)}, y_t^{(n)})) : (x_t^{(n)}, y_t^{(n)}) \in C_t\}$
 # Key Imagination
 $h_t^x \leftarrow \text{RNN}(\tilde{X}_{t-1}, h_{t-1}^x, r_t^c)$
 $\tilde{X}_t \leftarrow f_X(h_t^x)$
 # Value Imagination
 for $k \in [K]$ **in parallel do**
 $\tilde{h}_t^{(k)} \leftarrow \text{RNN}_k(\tilde{x}_t^{(k)}, \tilde{v}_{t-1}^{(k)}, \tilde{h}_{t-1}^{(k)})$
 end for
 $\{\tilde{v}_t^{(k)}\}_{k \in [K]} \leftarrow \text{Attend}(C_t \cup \{(\tilde{x}_t^{(k)}, \tilde{h}_t^{(k)})\}_{k \in [K]}; \tilde{X}_t)$
 # Store
 $\tilde{C}_t \leftarrow \{(\tilde{x}_t^{(k)}, \tilde{v}_t^{(k)})\}_{k \in [K]}$
end for

key-value pair $(\tilde{x}_t^{(k)}, \tilde{v}_t^{(k)})$ for $k \in [K] \equiv \{1, \dots, K\}$. We denote the imagined key-set by $\tilde{X}_t = \{\tilde{x}_t^{(k)}\}_k$ and the imagined value-set by $\tilde{V}_t = \{\tilde{v}_t^{(k)}\}_k$. When generating imaginary context, we also use the real context C_t gathered from the current task τ_t to inform the RMR about the characteristics of the current task. This is summarized by:

$$\tilde{C}_t, \tilde{\mathcal{H}}_t = \text{RMR}(C_t, \tilde{C}_{t-1}, \tilde{\mathcal{H}}_{t-1}) \quad (5)$$

where $\tilde{\mathcal{H}}_t = \tilde{\mathcal{H}}_t^{\text{key}} \cup \tilde{\mathcal{H}}_t^{\text{val}}$ is a set of RMR's hidden states, which consists of the hidden-states of key-generation and value-generation, respectively. RMR generates imaginary keys first and then, conditioning on the generated keys, the imaginary values.

3.1.1. GENERATING IMAGINARY KEYS

The main goal of imaginary key generation is to update the formation of the imaginary inputs (i.e., keys) upon a task-shift so that it can provide attention on a meaningful area of the new task. This requires the model to see (i) which input areas were tracked so far and to learn the task-shift dynamics via \tilde{X}_{t-1} and (ii) what the current task is via C_t . To this end, RMR deploys an RNN called *key-tracker*. Therefore, the key-generation process becomes:

$$(\tilde{X}_t, \tilde{h}_t^{\text{key}}) = \text{GenerateKey}(C_t, \tilde{X}_{t-1}, \tilde{h}_{t-1}^{\text{key}}). \quad (6)$$

where \tilde{h}_t^{key} is the hidden state of the key-tracker.

Specifically, we first concatenate the previous imaginary-keys $\tilde{x}_{t-1}^{(1)}, \dots, \tilde{x}_{t-1}^{(K)}$ to make the input representation and then encode the real context C_t using an order-invariant encoder $r_t = \sum_n \text{MLP}(x_t^{(n)}, y_t^{(n)})$. We then provide the concatenated keys and the real context encoding to the key-tracker RNN as the inputs. The new imaginary key-set \tilde{X}_t

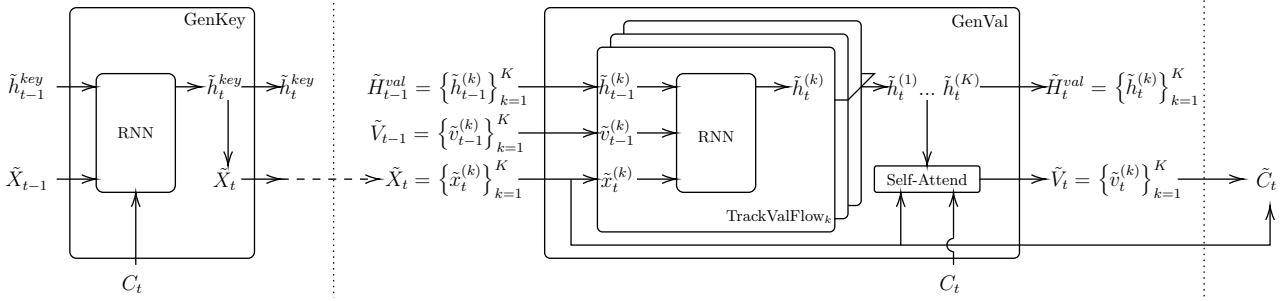


Figure 2: Illustration of Recurrent Memory Reconstruction (RMR) for generating imaginary context.

is generated from the updated hidden state \tilde{h}_t^{key} . To compute the order-invariant encoding, other implementation options are described in Garnelo et al. (2018); Eslami et al. (2018); Gordon et al. (2019).

3.1.2. GENERATING IMAGINARY VALUES

Given the new imaginary keys, we generate the corresponding imaginary values. Similar to the key generation process, to generate a new value set \tilde{V}_t , RMR takes the new key \tilde{X}_t , the new real context C_t , and the previous value-set \tilde{V}_{t-1} as inputs. This results in the following summary of the value-generation process:

$$\tilde{V}_t, \tilde{H}_t^{val} = \text{GenerateValue}(\tilde{X}_t, C_t, \tilde{V}_{t-1}, \tilde{H}_{t-1}^{val}) \quad (7)$$

where \tilde{H}_{t-1}^{val} is a set of hidden states for value generation.

In designing the value generator, two principles play key roles. First, value-generation process should be aware of what has happened in the past to generate a useful value for the current task upon a task-shift. Second, the values in a value-set should be aware of each other in order to obtain an optimal formation of the values. To realize this, we implement the value generation by the following two components: value-flow tracking and value-flow interaction.

i) Value-Flow Tracking. The purpose of this stage is to implement recurrence that captures value-transitions across the tasks. To this end, for each memory cell $k \in [K]$, we assign an RNN, TrackValueFlow updating:

$$\tilde{h}_t^{(k)} = \text{TrackValueFlow}(\tilde{x}_t^{(k)}, \tilde{v}_{t-1}^{(k)}, \tilde{h}_{t-1}^{(k)}) \quad (8)$$

where $\tilde{h}_t^{(k)}$ is the RNN hidden-state which we call the *value-tracker state*. Here, each value-tracker state acts as a proposal for the final imaginary value, and each RNN is seen as *tracking the flow* of the value in a particular memory cell, hence the name.

ii) Value-Flow Interaction. In the previous step, TrackValueFlow, each value was updated using the new key and its previous value. However, this step does not know what other values were generated. Hence, we require an interactive update. We revise each proposal value via self-attention on other values in the proposal value-set and the real context set.

Specifically, we first use the proposal values to construct a proposal key-value set $\tilde{C}_t^{\text{prop}} = \{(\tilde{x}_t^{(k)}, \tilde{h}_t^{(k)})\}_{k \in [K]}$. We then combine it with the real context C_t to get $C_t \cup \tilde{C}_t^{\text{prop}}$. Finally, we perform self-attention on this union using the imaginary keys as the attention queries and obtain the final imaginary value-set $\tilde{V}_t = \{\tilde{v}_t^{(k)}\}_{k \in [K]}$:

$$\tilde{V}_t = \text{Self-Attend}(C_t \cup \tilde{C}_t^{\text{prop}}; \tilde{X}_t). \quad (9)$$

This completes the generation of imaginary context, $\tilde{C}_t = (\tilde{X}_t, \tilde{V}_t)$. Algorithm 1 and Fig. 2 show the described process of generating imaginary context.

3.1.3. READING IN RMR

To perform a read operation on the RMR at a given task-step t , we propose to perform attention on the extended context $\tilde{C}_t = C_t \cup \tilde{C}_t$. Given a query input x_t^q , we thus obtain the read value as follows $r_{x_t^q} = \text{Attend}(\tilde{C}_t; x_t^q)$.

3.2. Attentive Sequential Neural Processes

Using the imagined context obtained from RMR, we can now address the problems of sparse and obsolete context in SNP. In this section, we describe how we augment Sequential Neural Processes with RMR and, in doing so, propose *Attentive Sequential Neural Processes-RMR*.

3.2.1. GENERATIVE PROCESS

We equip the observation model in SNP with the augmented memory \tilde{C}_t provided by RMR. With this, the observation model becomes $P(y_t|x_t, z_t, \tilde{C}_t)$ (see Fig. 3). Then, given a target input x_t , the observation model reads \tilde{C}_t to obtain an attention encoding $r_{x_t} = \text{Attend}(\tilde{C}_t; x_t)$.

In this way, combining SNP with RMR, the generative process is described as follows: $P(Y, Z, \tilde{C}|X, C) =$

$$\prod_{t=1}^T \underbrace{P(y_t|x_t, z_t, \tilde{C}_t)P(z_t|z_{<t}, \tilde{C}_t)}_{\text{SNP with access to RMR}} \underbrace{P(\tilde{C}_t|\tilde{C}_{<t}, C_{\leq t})}_{\text{RMR Update}}. \quad (10)$$

We call this model Attentive Sequential Neural Process with RMR (ASNP-RMR).

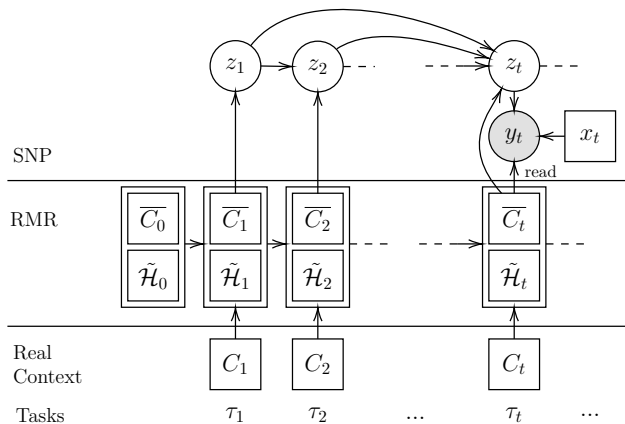


Figure 3: Illustration of the generative process of ASNP-RMR.

3.2.2. LEARNING AND INFERENCE

Because the true posterior is not tractable, ASNP is trained using a variational approximation with the following autoregressive formulation:

$$P(Z|\tilde{C}, C, D) \approx \prod_{t=1}^T Q(z_t|z_{<t}, \tilde{C}_{\leq t}, C, D) \quad (11)$$

where $D = (X, Y)$ and C is the set aggregations of C_t over the entire roll-out. For training, the following ELBO is maximized w.r.t. θ and ϕ : $\mathcal{L}_{\text{ASNP}}(\theta, \phi) =$

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{Q_{\phi}(z_t|C, D)} \left[\log P_{\theta}(y_t|x_t, z_t, \tilde{C}_t, C_t) \right] \\ & - \mathbb{E}_{Q_{\phi}(z_{<t})} \left[\mathbb{KL}(Q_{\phi}(z_t|z_{<t}, \tilde{C}_{\leq t}, C, D) \parallel \right. \\ & \quad \left. P_{\theta}(z_t|z_{<t}, \tilde{C}_{\leq t}, C_{\leq t})) \right]. \end{aligned}$$

For backpropagation, we use the reparameterization trick (Kingma & Welling, 2013). See Appendix C for derivation.

4. Related Works

Meta-learning approaches have become attractive for learning to learn new tasks at test-time. In this line of work, GQN (Eslami et al., 2018) renders 3D scenes from a few viewpoints, and NP (Garnelo et al., 2018) generalizes GQN. Subsequently, ANP (Kim et al., 2019) identifies and resolves the problem of underfitting in NP by using query-dependent representation. Similarly, in Rosenbaum et al. (2018), the authors introduce attention to GQN to render complex 3D scenes in large procedurally-generated maps as in Minecraft. Functional Neural Processes (FNP) (Louizos et al., 2019) learns a graph of dependencies between a pre-selected set of points and the training points for modeling distributions over functions. ANP-RNN (Qin et al., 2019) encodes the target inputs via an RNN and use the hidden states as queries

to attend on the context points in a meta-learning setting. To extend NP to *meta-transfer learning*, SNP (Singh et al., 2019) models a sequence of tasks with sequential latent representations. It outperforms NP in meta-transfer learning a sequence of tasks that come from different but related distributions. Recurrent Neural Process (RNP) (Kumar, 2019) also deals with meta-transfer learning, but it transfers via deterministic representations and shows high uncertainty with sparse context. Willi et al. (2019) also propose a model named RNP for meta-transfer learning.

The term *meta-transfer learning* has also been used in connection to a different set of problems that deal with discovery of causal mechanisms (Bengio et al., 2019), fast-adaption from a large-scale trained model (Sun et al., 2019), and knowledge transfer to different architectures and tasks (Jang et al., 2019). In Kang & Feng (2018), the authors propose transferable meta-learning to apply a meta-trained model to a task from a task distribution different from the trained.

Santoro et al. (2018) and Goyal et al. (2019) propose approaches combining self-attention and recurrent update. Recurrent Memory Core (RMC) (Santoro et al., 2018) self-attends the entire memory with a single input vector and updates the memory through an RNN module. Recurrent Independent Mechanisms (RIM) (Goyal et al., 2019) recurrently updates the memory but self-attends only the memories that are estimated to be the most relevant using the visual input. Although these methods combine recurrent and attention modules, however unlike RMR, their memory elements and inputs contain only values. They do not deal with key-value pairs in a meta-transfer learning setting.

5. Experiments

In this section, we describe our experiments to answer two key questions: *i)* By resolving the problems of sparse context and obsolete context, can we improve meta-transfer learning? *ii)* If yes, what are the needed memory sizes and computational overhead during training? We also perform an ablation on RMR to demonstrate the need for flow-tracking and flow-interaction. In the rest, we first describe the baselines and our experiment settings. We then describe our results on dynamic 1D regression, dynamic 2D image completion, and dynamic 2D image rendering.

	ANP	SNP	ASNP-W	ASNP-RMR
Sequential Latent	✗	✓	✓	✓
Attention	✓	✗	✓	✓
Recurrent Memory	✗	✗	✗	✓

Table 1: Taxonomy of the models considered for evaluation.

Baselines. We consider three baselines – ANP, SNP and ASNP-W. These are characterized by whether or not they contain a sequential latent, attention mechanism or a recurrent memory (see Table 1). Among these, ASNP-W is an




Figure 4: *Top:* Target-NLL computed at each time-step in the task sequence. *Bottom:* Target-NLL convergence against the wall-clock time computed during training on a held-out set. Hidden unit size in all models is 128.

extension of SNP such that the observation model attends on a window of K -most recent contexts. Hence, ASNP-W contains sequential latent and attention but *not* a recurrent memory. These baselines are to test whether sequential latent and the standard attention alone can solve underfitting without addressing the problems of sparse and obsolete context. We also test ANP and SNP using different latent sizes to investigate its effect on performance. Similarly, we test ASNP-RMR and ASNP-W using different memory sizes and analyze its effect. See Appendix A for more details on their implementation.

Performance Metric and Context Regimes. We evaluate our models on hundreds of held-out sequences of tasks and analyze their performance in modeling the target outputs. Our performance metric is *Target-NLL* defined as $-\mathbb{E}_{z \sim P(Z|C)} \log P(Y|X, Z, C)$. For dynamic 1D regression and dynamic 2D image completion, we consider two regimes for evaluation – *i) Sparse-Context Regime.* In this regime, we consider task-sequences of length 50. We provide a sparse context for 45 randomly chosen tasks and empty context for the remaining. This regime tests the model’s ability to transfer-learn from sparse contexts gathered from *different* tasks. At every increment of the task-step, we expect the model’s performance to improve as the model collects more context. *ii) Transfer-Prediction Regime.* In this regime, we consider task-sequences of length 20. We provide a large-sized context for the first 10 task-steps and empty context for the remaining. This regime allows the model to infer the first 10 tasks and their dynamics with high certainty. Subsequently, the model must use this information to transfer-learn the remaining 10 tasks. At every increment of the task-step, we expect the performance of any model to deteriorate as the contexts become more obsolete. However, if a model can deal with this problem, we




Figure 5: Samples for dynamic 1D regression in sparse-context regime at $t = 33$. *Dotted line:* True function. *Blue line:* Predicted function. *Shaded light-blue region:* Uncertainty. *Black dot:* Context points at $t = 33$. *Blue dots:* Past context points. Darker dots are more recent. *Dashed-rectangles:* Regions where the predictions are inaccurate.

expect it to show less degradation.

For 2D dynamic image rendering also, we experiment on the sparse-context and transfer-prediction regimes but with task-sequences of length 6. See Appendix B for details about sequence lengths and context sizes.

5.1. Dynamic 1D Regression

In this setting, the tasks are 1D functions that change with time. To generate this dataset, we draw a function from a GP at each task-step such that the kernel-parameters of the GP change according to some linear dynamics. Hence, the model must estimate the function at unseen points and also track the shifts in its shape. We used a squared-exponential kernel for GPs. See Appendix B.1 for more details.

Target-NLL. We plot the target-NLLs for each task-step in Fig. 4. Consider the performance of ASNP-W to study the effect of extending SNP with the standard attention. We observe that improvements compared to SNP are small. Comparing these with ASNP-RMR shows that sequential latent and the standard attention alone cannot solve underfitting without addressing sparse and obsolete contexts. In *transfer-prediction* regime, for $t \leq 10$, we note that target-NLLs of ANP are similar to that of ASNP-RMR. This is because, in this interval, ANP exploits the larger context sizes without attending to the obsolete points. However, the main focus of this setting is multi-step transfer for $t > 10$ on which ANP degrades quickly.

Effect of Memory Size and Latent Size. In Fig. 6, we show the effect of memory and latent sizes on performance. We observe that increasing the latent size in SNP from 128 to 1024 shows small improvements with diminishing returns. This shows that a larger latent size is not adequate. Additionally, we observe that ASNP-W does not surpass ANP significantly. Hence, we conclude that the recurrent




Figure 6: Target-NLL computed as a function of memory sizes in sparse-context regime. Here, h denotes the latent size.




Figure 7: Target-NLL computed at each time-step in moving CelebA image completion task sequence. Hidden unit size in all models is 128.

latents $z_{\leq t}$ fail to inform the observation model about the shift in the obsolete points. We further note that ASNP-RMR at $K = 9$ outperforms ASNP-W for all choices of K . This shows that RMR is more size-efficient.

Training Time. In Fig. 4, we plot the target-NLL against the training wall-clock time, and we note that ASNP-RMR imposes no significant overhead in convergence.

Qualitative Analysis. In Fig. 5, we show the predictive means of the target function. We observe that SNP underfits the target function in multiple locations. Furthermore, in ANP and ASNP-W, attention helps the prediction only at those obsolete points that did not undergo a significant shift. In these models, attention also causes misestimation of the function when obsolete points shift significantly. It shows that without addressing the obsolete context issue, the standard attention cannot resolve underfitting. In comparison, ASNP-RMR makes more accurate predictions.

5.2. Dynamic 2D Image Completion

In this setting, the task is to complete an image by estimating the pixel value at a given pixel location. The image belongs to a sequence of images that contains a moving image patch on a white canvas. In this dynamic setting, the model must not only estimate the unseen pixel values but also track their

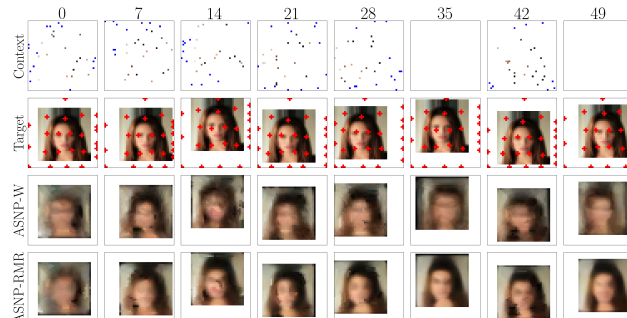


Figure 8: Qualitative evaluation on moving CelebA image completion task sequence. *Red dots*: Imaginary keys imagined by ASNP-RMR. The memory size K is 25.

motion. The moving images are taken from the MNIST (LeCun et al., 1998) and the CelebA (Liu et al., 2015) datasets, and hence we call these settings *moving MNIST* and *moving CelebA*, respectively.

Target NLL. We plot the target-NLLs in Fig. 7. Recall that in the sparse-context regime, a model must transfer-learn at every task-step, and in the transfer-prediction regime, a model must transfer-learn on task-steps $t > 10$ that have an empty context. We observe that when such transfer-learning is required, the performance of ASNP-W degrades compared to SNP. This implies that attention on obsolete points is not only ineffective but also detrimental.

Effect of Memory and Latent Sizes. From Fig. 6, our conclusions about the effect of memory size and latent size are similar to those in dynamic 1D regression. ASNP-RMR is not only size-efficient but also shows monotonic improvement as we increase the memory size. Note that in Fig. 6, we do not test ANP using a larger latent size because we find its memory usage too large to train.

Qualitative Analysis. Fig. 8 shows the positions of imaginary keys generated by RMR. We observe that the model learns to discover certain key points on the CelebA patch, and it tracks them as the patch moves. It also places some keys on the edges of the canvas to account for the bounces. Fig. 9 shows qualitative samples for moving MNIST image completion in sparse-context setting. As ASNP-RMR accumulates more context, RMR reforms the obsolete points for the current task. Consequently, we observe that the predicted images become clearer over time. We also observe that in comparison to SNP, ASNP-W shows deterioration.




Figure 9: Moving MNIST image completion samples in the sparse-context regime. Memory size of ASNP-RMR is 25.

In Fig. 9 and 8, consider the task-steps with an empty context. On these task-steps, ANP and ASNP-W predict poor quality images compared to ASNP-RMR, which shows that the standard attention cannot deal with the obsolete context.

5.3. Ablation Study on RMR

In this section, we perform an ablation on RMR to create two modifications, and we report the results in Fig. 10.

i) No Value-Flow Tracking. In this modification, the model generates new imaginary values only via self-attention on the previous values. Thus, without the recurrence, we expect the model to forget the transition dynamics. In Fig. 10, we observe in the transfer-prediction setting that the performance deteriorates sharply in task-steps [11, 20]. This is because the context is empty, and value imagination fails to extrapolate without capturing the transition dynamics.

ii) No Value-Flow Interaction. In this modification, the model generates new imaginary values only via Value-Flow Tracking. To incorporate the real context, we provide it as an input to the flow tracking RNNs. Because of no value interaction, we observe performance degradation in Fig. 10. Thus, we conclude that to perform effective value imagination, the model requires both flow tracking and interaction.

5.4. Dynamic 2D Image Rendering

In this setting, we consider a sequence of images as in the moving CelebA dataset. The task for the model is to take a location on the canvas and predict the image-patch centered at the location. Because the task is to generate an image, we replace the SNP baseline with a TGQN (Singh et al., 2019). Similarly, we replace ASNP-RMR and ASNP-W with ATGQN-RMR and ATGQN-W, respectively. Here, ATGQN stands for attentive TGQN – a model that we develop which incorporates attention into TGQN. Hence, ATGQN-RMR is TGQN equipped with attention on RMR,

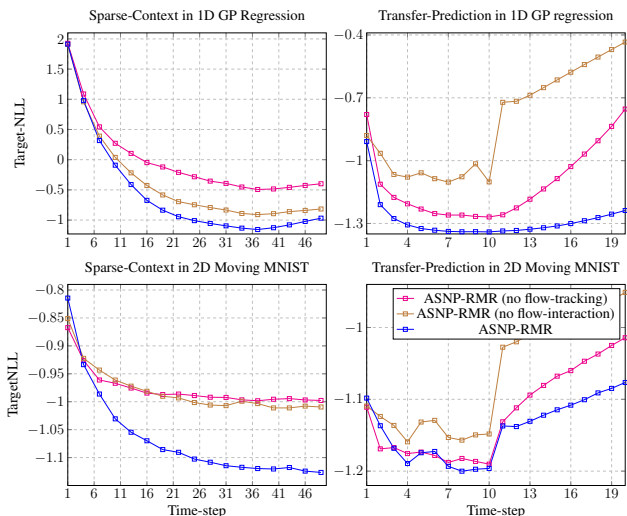


Figure 10: Target-NLL for ASNP-RMR compared against two alternate versions by way of ablation – *i)* without flow-tracking and *ii)* without flow-interaction.

Regime	ATGQN -RMR	ATGQN -W	TGQN
Sparse-Context	2.57	2.77	3.23
Transfer-Prediction	2.04	2.8	3.08

Table 2: Target-NLL comparison between ATGQN-RMR, ATGQN-W and TGQN in moving CelebA image rendering task.

and ATGQN-W is equipped with the standard attention on a context window. See Appendix B.3 for implementation details. We report the results on these models in Table 2, and we note that ATGQN-RMR outperforms ATGQN-W and TGQN.

6. Conclusion

In this paper, we argued that the two problems, sparse context and obsolete context, observed in meta-transfer learning due to temporal task-shift, make the underfitting issue in SNP more severe. Then, to resolve this, we proposed a novel attention model using imaginary context generated by Recurrent Memory Reconstruction (RMR) and a robust probabilistic meta-transfer learning model, Attentive Sequential Neural Processes. Our experiments demonstrate that existing methods show weaknesses in dealing with sparse and obsolete context, and that using RMR-based attention in SNP is an effective way to resolve the issue. The ablation study shows that the recurrent context modeling and the interaction model are the key components to achieve this improved robustness. In the future, it will be interesting to apply this robust model to various other problems including meta-transfer reinforcement learning.

Acknowledgements

JY thanks Kakao Brain and SAP for their support and Hoyong Jang for helpful feedback. SA thanks Kakao Brain and Center for Super Intelligence (CSI) for their support.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Bengio, Y., Bengio, S., and Cloutier, J. *Learning a synaptic learning rule*. Université de Montréal, Département d’informatique et de recherche . . . , 1990.
- Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., and Pal, C. A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*, 2019.
- Eslami, S. A., Rezende, D. J., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., Ruderman, A., Rusu, A. A., Danihelka, I., Gregor, K., et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.
- Gordon, J., Bruinsma, W. P., Foong, A. Y., Requeima, J., Dubois, Y., and Turner, R. E. Convolutional conditional neural processes. *arXiv preprint arXiv:1910.13556*, 2019.
- Goyal, A., Lamb, A., Hoffmann, J., Sodhani, S., Levine, S., Bengio, Y., and Schölkopf, B. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*, 2019.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*, 2018.
- Jang, Y., Lee, H., Hwang, S. J., and Shin, J. Learning what and where to transfer. *arXiv preprint arXiv:1905.05901*, 2019.
- Kang, B. and Feng, J. Transferable meta learning across domains. In *UAI*, pp. 177–187, 2018.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kumar, A., Eslami, S., Rezende, D. J., Garnelo, M., Viola, F., Lockhart, E., and Shanahan, M. Consistent generative query networks. *arXiv preprint arXiv:1807.02033*, 2018.
- Kumar, S. *Spatiotemporal Modeling using Recurrent Neural Processes*. PhD thesis, Carnegie Mellon University Pittsburgh, PA, 2019.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Louizos, C., Shi, X., Schutte, K., and Welling, M. The functional neural process. *arXiv preprint arXiv:1906.08324*, 2019.
- Nair, V. and Hinton, G. E. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807–814, 2010.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10): 1345–1359, 2009.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.
- Pratt, L. Y. Discriminability-based transfer between neural networks. In *Advances in neural information processing systems*, pp. 204–211, 1993.
- Qin, S., Zhu, J., Qin, J., Wang, W., and Zhao, D. Recurrent attentive neural process for sequential data. *arXiv preprint arXiv:1910.09323*, 2019.
- Rosenbaum, D., Besse, F., Viola, F., Rezende, D. J., and Eslami, S. Learning models for visual 3d localization with implicit mapping. *arXiv preprint arXiv:1807.03149*, 2018.
- Santoro, A., Faulkner, R., Raposo, D., Rae, J., Chrzanowski, M., Weber, T., Wierstra, D., Vinyals, O., Pascanu, R., and Lillicrap, T. Relational recurrent neural networks. In *Advances in neural information processing systems*, pp. 7299–7310, 2018.
- Schmidhuber, J. *Evolutionary Principles in Self-referential Learning: On Learning how to Learn: the Meta-meta-meta...-hook*. 1987.
- Singh, G., Yoon, J., Son, Y., and Ahn, S. Sequential neural processes. *arXiv preprint arXiv:1906.10264*, 2019.

- Sun, Q., Liu, Y., Chua, T.-S., and Schiele, B. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 403–412, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wagstaff, E., Fuchs, F. B., Engelcke, M., Posner, I., and Osborne, M. On the limitations of representing functions on sets. *arXiv preprint arXiv:1901.09006*, 2019.
- Willi, T., Masci, J., Schmidhuber, J., and Osendorfer, C. Recurrent neural processes. *arXiv preprint arXiv:1906.05915*, 2019.
- Xingjian, S., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., and Woo, W.-c. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pp. 802–810, 2015.